

# **Comparison of QNX Neutrino, Windows CE7, and Linux RT operating systems on ATOM processor**

---

## **Copyright**

© Copyright Dedicated Systems Experts NV. All rights reserved, no part of the contents of this document may be reproduced or transmitted in any form or by any means without the written permission of Dedicated Systems Experts NV, Diepenbeemd 5, B-1650 Beersel, Belgium.

## **Disclaimer**

Although all care has been taken to obtain correct information and accurate test results, Dedicated Systems Experts, VUB-Brussels, RMA-Brussels and the authors cannot be liable for any incidental or consequential damages (including damages for loss of business, profits or the like) arising out of the use of the information provided in this report, even if these organizations and authors have been advised of the possibility of such damages.

---

## **Authors**

Luc Perneel (1, 2), Hasan Fayyad-Kazan(2) and Martin Timmerman (1, 2, 3)  
1: Dedicated Systems Experts, 2: VUB-Brussels, 3: RMA-Brussels

---

<http://download.dedicated-systems.com>

E-mail: [info@dedicated-systems.com](mailto:info@dedicated-systems.com)

## EVALUATION REPORT LICENSE

This is a legal agreement between you (the downloader of this document) and/or your company and the company DEDICATED SYSTEMS EXPERTS NV, Diepenbeemd 5, B-1650 Beersel, Belgium.

It is not possible to download this document without registering and accepting this agreement on-line.

1. **GRANT.** Subject to the provisions contained herein, Dedicated Systems Experts hereby grants you a non-exclusive license to use its accompanying proprietary evaluation report for projects where you or your company are involved as major contractor or subcontractor. You are not entitled to support or telephone assistance in connection with this license.
2. **PRODUCT.** Dedicated Systems Experts shall furnish the evaluation report to you electronically via Internet. This license does not grant you any right to any enhancement or update to the document.
3. **TITLE.** Title, ownership rights, and intellectual property rights in and to the document shall remain in Dedicated Systems Experts and/or its suppliers or evaluated product manufacturers. The copyright laws of Belgium and all international copyright treaties protect the documents.
4. **CONTENT.** Title, ownership rights, and an intellectual property right in and to the content accessed through the document is the property of the applicable content owner and may be protected by applicable copyright or other law. This License gives you no rights to such content.
5. **YOU CANNOT:**
  - You cannot, make (or allow anyone else make) copies, whether digital, printed, photographic or others, except for backup reasons. The number of copies should be limited to 2. The copies should be exact replicates of the original (in paper or electronic format) with all copyright notices and logos.
  - You cannot, place (or allow anyone else place) the evaluation report on an electronic board or other form of on line service without authorisation.
6. **INDEMNIFICATION.** You agree to indemnify and hold harmless Dedicated Systems Experts against any damages or liability of any kind arising from any use of this product other than the permitted uses specified in this agreement.
7. **DISCLAIMER OF WARRANTY.** All documents published by Dedicated Systems Experts on the World Wide Web Server or by any other means are provided "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. This disclaimer of warranty constitutes an essential part of the agreement.
8. **LIMITATION OF LIABILITY.** Neither Dedicated Systems Experts nor any of its directors, employees, partners or agents shall, under any circumstances, be liable to any person for any special, incidental, indirect or consequential damages, including, without limitation, damages resulting from use of OR RELIANCE ON the INFORMATION presented, loss of profits or revenues or costs of replacement goods, even if informed in advance of the possibility of such damages.
9. **ACCURACY OF INFORMATION.** Every effort has been made to ensure the accuracy of the information presented herein. However Dedicated Systems Experts assumes no responsibility for the accuracy of the information. Product information is subject to change without notice. Changes, if any, will be incorporated in new editions of these publications. Dedicated Systems Experts may make improvements and/or changes in the products and/or the programs described in these publications at any time without notice. Mention of non-Dedicated Systems Experts products or services is for information purposes only and constitutes neither an endorsement nor a recommendation.
10. **JURISDICTION.** In case of any problems, the court of BRUSSELS-BELGIUM will have exclusive jurisdiction.

**Agreed by downloading the document via the internet.**

## Contents

1	About the RTOS evaluation project .....	5
1.1	Purpose and scope of the RTOS evaluation.....	5
1.2	Test framework used: 2.9.....	5
2	About the OSs and the testing platform.....	6
2.1	Software .....	6
2.2	Hardware.....	7
3	Evaluation results overview .....	8
3.1	Dedicated Systems' ratings for the tested RTOSs.....	8
3.2	Rating Criteria .....	9
3.3	Positive and negative points for each OS .....	10
3.4	Ratings by category .....	11
3.4.1	QNX Neutrino 6.5 with Patch 2530.....	12
3.4.2	Windows Embedded Compact 7 .....	12
3.4.3	Linux 2.6.33.7 with rt30 Patch .....	12
3.5	Tests Summary .....	13
3.5.1	Clock tick processing duration (CLK-P-DUR) .....	13
3.5.2	Thread switch latency between same priority threads (THR-P-SLS).....	14
3.5.3	Maximum sustained interrupt frequency (IRQ_S_SUS).....	15
3.5.4	Mutex acquire-release timings: contention case (MUT-P-ARC).....	16
3.5.5	Mutex acquire-release timings: no-contention case (MUT-P-ARN) .....	17
4	Detailed Comparison .....	18
4.1	Clock tests (CLK) .....	18
4.2	Thread tests (THR).....	19
4.2.1	Thread creation behaviour (THR-B-NEW).....	19
4.2.2	Round robin behaviour (THR-B-RR) .....	20
4.2.3	Thread switch latency between same priority threads (THR-P-SLS).....	21
4.2.4	Thread creation and deletion time (THR-P-NEW).....	23
4.3	Semaphore tests (SEM).....	25
4.3.1	Semaphore locking test mechanism (SEM-B-LCK) .....	25
4.3.2	Semaphore releasing mechanism (SEM-B-REL) .....	26
4.3.3	Time needed to create and delete a semaphore (SEM-P-NEW) .....	26
4.3.4	Test acquire-release timings: non-contention case (SEM-P-ARN) .....	28
4.3.5	Test acquire-release timings: contention case (SEM-P-ARC) .....	29
4.4	Mutex tests (MUT).....	31
4.4.1	Priority inversion avoidance mechanism (MUT-B-ARC) .....	31
4.4.2	Mutex acquire-release timings: no-contention case (MUT-P-ARN) .....	32
4.4.3	Mutex acquire-release timings: contention case (MUT-P-ARC).....	33
4.5	Interrupt tests (IRQ).....	35
4.5.1	Interrupt latency (IRQ_P_LAT) .....	35
4.5.2	Interrupt dispatch latency (IRQ_P_DLT) .....	37

Doc: **EVA-2.9-CMP-ATOM**

Issue: **v 1.00**

Date: **May 1, 2012**

4.5.3	Interrupt to thread latency (IRQ_P_TLT) .....	38
4.5.4	Maximum sustained interrupt frequency (IRQ_S_SUS) .....	39
5	Conclusion .....	40
6	Related documents .....	41
7	Appendix A: Acronyms .....	42

## 1 About the RTOS evaluation project

This section describes the purpose and scope of the evaluations conducted by Dedicated Systems.

### 1.1 Purpose and scope of the RTOS evaluation

This document provides quantitative measures to help potential RTOS users make objective comparisons between OSs and help them decide which OS is better for their needs.

This document compares the results of the quantitative evaluations of three real time operating systems (RTOSs). These OSs are:

- QNX Neutrino 6.5 patch 2530
- Windows Embedded Compact 7
- Linux 2.6.33.7.2-rt30

The order in which we list the OSs is based on the overall results obtained by the OSs, with the OS with the best results listed first and the others following in descending order. This ordering is maintained throughout the whole report.

These RTOSs were evaluated on the same ATOM platform (Advantech SOM-6760).

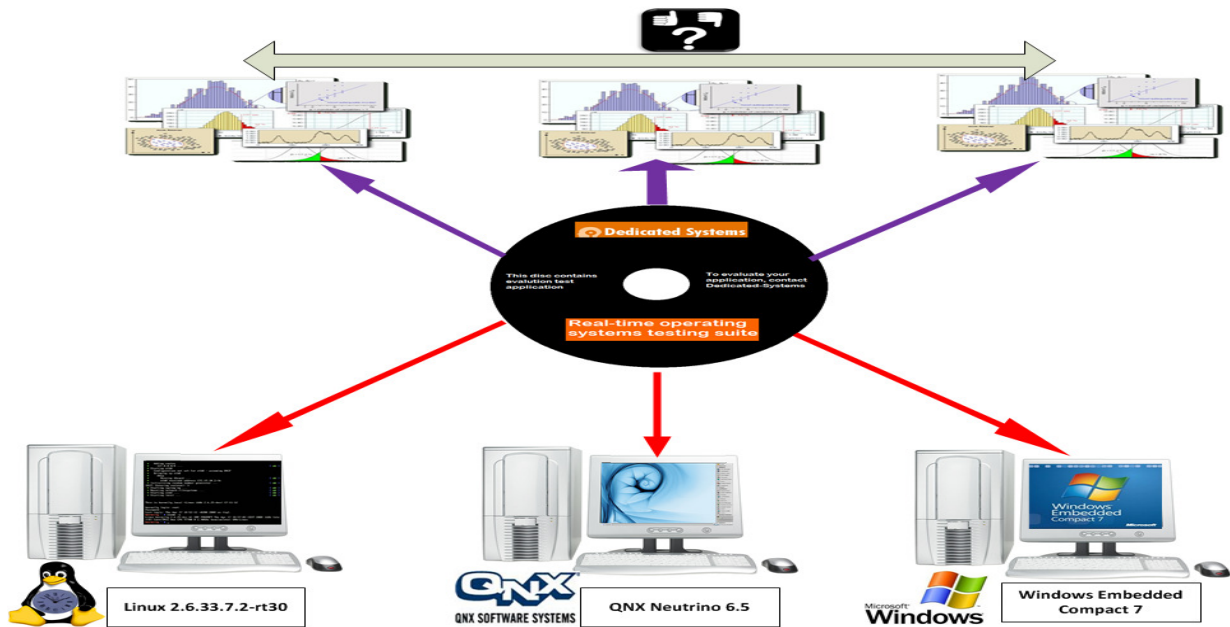


Figure 1: High level view of the evaluation procedure

### 1.2 Test framework used: 2.9

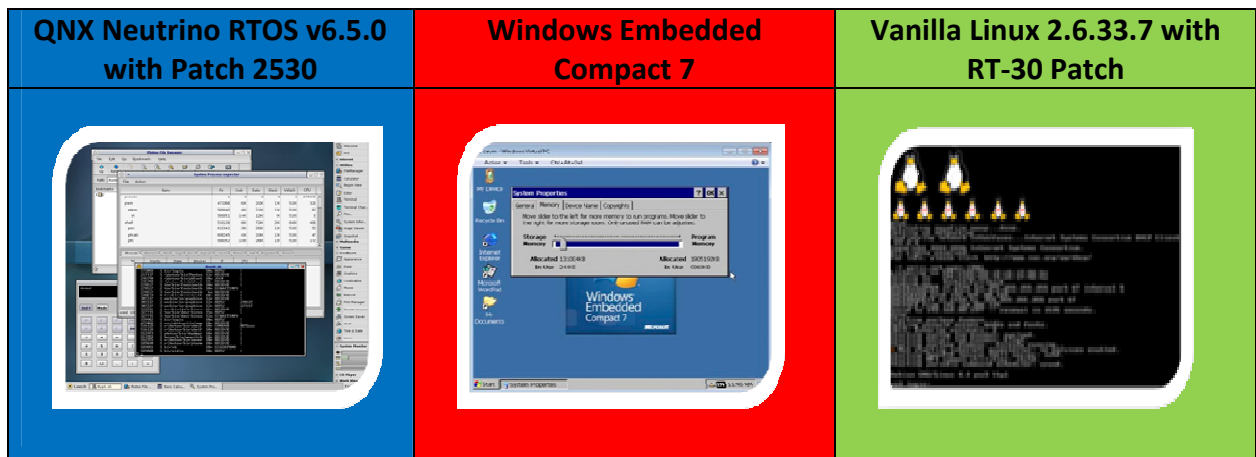
This document shows the test results in the scope of the evaluation framework 2.9. More details about this framework are found in Doc 1 (see section 6).

## 2 About the OSs and the testing platform

This section describes the OSs that Dedicated Systems tested using its Evaluation Testing Suite, and the hardware on which these OSs were running during the testing.

### 2.1 Software

The following table shows the operation systems' versions whose behavior and performance results were compared by Dedicated Systems after testing them with its evaluation testing suite on the same ATOM platform (Advantech SOM-6760).



**Table 1: The evaluated OSs**

For **QNX Neutrino 6.5**, Patch 2530 was applied. This patch introduces a fix to the io-pkt network stack where a timer pulse implementation is used instead of attaching a handler to the timer interrupt. This patch significantly improves clock tick processing times and results in improved real time performance.

For **Windows Embedded Compact 7**, no patches were applied.

For “**Vanilla**” **Linux 2.6.33.7**, real-time patch rt-30 was applied to provide some real time characteristics for the Linux kernel. This RT patch was the latest version officially released by OSADL.

## 2.2 Hardware

We conducted our tests on the same ATOM platform. This platform has the following characteristics:

- Motherboard: Advantech SOM-6760, PCI bus at 33MHz, using the System Controller Hub US15W.
- CPU: Intel Atom Z530 1.6GHz 133MHz Front Side Bus.
- 32KByte L1 Instruction Cache,
- 24KByte L1 Write Back Data Cache,
- 512KByte 8-way L2 Cache (which can be reduced up to zero in some processor sleep states)
- 1 core with hyper-threading support (however hyper threading was disabled during this test).
- RAM: 512MB DDR2
- VMETRO PCI exerciser in PCI slot 3 (PCI interrupt level D, local bus interrupt level 10)
- VMETRO PBT-315 PCI analyser in PCI slot 4.
- External and CPU internal cache was enabled during the tests.




Figure 2: The ATOM board on which the tests were conducted



## 3 Evaluation results overview

This section presents the overall ratings and evaluations based on key tests.

### 3.1 Dedicated Systems' ratings for the tested RTOSs

Table 2 shows Dedicated Systems' overall ratings for the tested OSs:

QNX Neutrino 6.5.0	Windows Embedded Compact 7	Linux 2.6.33.7-rt30
★★★★★	★★★★☆	★★★

**Table 2: Overall ratings for the evaluated OSs**

Both, QNX Neutrino 6.5 and Windows Embedded Compact 7 are excellent real-time operating systems. Windows Embedded Compact 7 got half a star less than QNX Neutrino 6.5 because its testing results are to some extent slower than the testing results of the QNX Neutrino 6.5. However, in practice, this difference will be too small to notice.

Although the RT patch was applied to Linux, it is still considered to be at a far distance from the other two traditional OSes. Further, the correct kernel configuration (at build and at run time) is required to achieve these results, something which is not that easy to achieve. Therefore we give Linux RT three (3) stars.





## 3.2 Rating Criteria

After testing each OS using the Dedicated Systems Evaluation Testing Suite, we used a star system to give each OS a rating based on the performance and behavior results. The maximum number of stars that an OS can receive is five (5).

Rating	Availability of real-time requirements	Performance, behavior and interrupts testing results
★★★★★	The OS works correctly out-of-the-box. No fear about the kernel configuration	Excellent (Hard RT is met)
★★★★	The OS works correctly out-of-the-box. No fear about the kernel configuration	Very good (Hard RT is met)
★★★	Special attention and knowledge are required to correctly configure the kernel	Good (only soft RT is met)
★★	Special attention and deep knowledge are required to correctly configure the kernel	Bad (even soft RT problems)
★	Correctly configuring the kernel is problematic	Problematic
☹	NO RT capabilities	Fails (for hard and soft RT applications)

**Table 3: Criteria used to evaluate OSs**

## 3.3 Positive and negative points for each OS

Evaluated OS	Positive points 	Negative points 
<b>QNX Neutrino 6.5.0</b>	1) Excellent architecture for a robust and distributed system. 2) Very fast and predictable performance. 3) Large number of (BSPs) and drivers can be easily downloaded. 4) The availability of documentation and support is very high. 5) Efficient and user friendly (IDE)	1) Not all code is available in source code. Customers can apply for source access.
<b>Windows Embedded Compact 7</b>	1) All protection primitives use priority inheritance, which is a major plus for achieving real-time behavior. 2) Good debugging tools which are also available for kernel/driver debugging. 3) Very easy to install and to set-up a target (from templates). 4) provides the same flexibility as a 32-bit general purpose OS	1) A lot of background information is only available for CE6R3. 2) Customizing the kernel and adding custom drivers (BSP) stays a daunting task once you go away from the default configurations.
<b>Linux 2.6.33.7-rt30</b>	1) No license fees. 2) Source code available. 3) Extensible	1) The real-time characteristics are present only when everything is configured and built correctly. 2) GPL is not completely free... 3) Setting up a complete embedded target from scratch is a daunting task.

**Table 4: Positive and negative points found in each evaluated OS**



## ***3.4.1 QNX Neutrino 6.5 with Patch 2530***

QNX Neutrino stands out as a clearly superior real-time OS compared to the other evaluated OSs. In addition to its design, which is much more robust and very easy to debug, even at the driver level, the data from our tests for this OS confirm that its real-time behaviour is considerably better than that of the other OSs. Further, this OS it is very well documented, and users do not have to worry about kernel configuration, as the OS kernel is always configured correctly.

## ***3.4.2 Windows Embedded Compact 7***

Like the QNX Neutrino RTOS, Microsoft Windows Compact Embedded 7 (CE7) is also a system built for real-time behavior. This OS has a large toolset, and many options are available for building different types of systems. However, this OS is a marginally slower than the QNX Neutrino 6.5, and the toolset for building drivers and develop BSPs is less intuitive than the ones supplied by QNX. Also like QNX Neutrino OS, it tested significantly better than Linux 2.6.33.7-rt30.

## ***3.4.3 Linux 2.6.33.7 with rt30 Patch***

The chief advantage of Linux is its open source licensing (no run-time fees). Note, however, that the GPL is not completely free, and investment is required to build a marketable system. For instance, though demo systems can be built quickly with Linux, the debugging, tuning and verification required to build a stable system ready for long-term use is much more difficult. Projects using Linux OSs tend to require large development teams. Further, projects that brew their own Linux flavor will need kernel experts who understand, for a start, how to set the kernel configurations (both at build and at run-time) to obtain real-time behavior.

## 3.5 Tests Summary

This section presents a brief comparative summary of the most important evaluation tests performed on the OSs we tested.

Detailed comparisons can be found in the next chapter. More detailed information about each test and its importance can be found in the corresponding documents: QNX Neutrino 6.5 [Doc 5], Windows Compact Embedded 7 (CE7) [Doc 9], and Linux 2.6.33.7-rt30 [Doc 7].

Note that in the comparison figures and tables:

- The lower values means better quality
- Values in the charts are in microseconds ( $\mu$ s)

### 3.5.1 Clock tick processing duration (CLK-P-DUR)

The “clock tick processing duration” test examines the clock tick processing duration in the kernel. The clock tick processing time is important because it impacts latencies everywhere in the system. The test results are extremely important because the clock interrupt will affect all the other measurements performed.




Figure 3a: Average clock interrupt duration

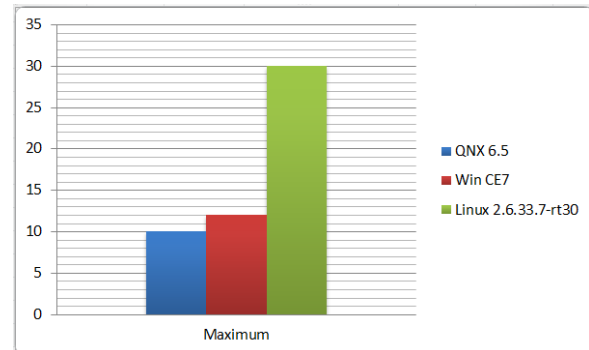



Figure 3b: Maximum clock interrupt duration

Since we are interested in real-time behavior, the maximum values are more important than the average values. From our results, it is clear that the traditional RTOSs (QNX and Windows CE) are still miles ahead of Linux. The maximum clock durations (Figure 3b) for QNX Neutrino 6.5 and Windows CE7 are almost the same (10 for QNX Neutrino 6.5 and 12 for Windows CE7), while it is very high (30) for the Linux RT.

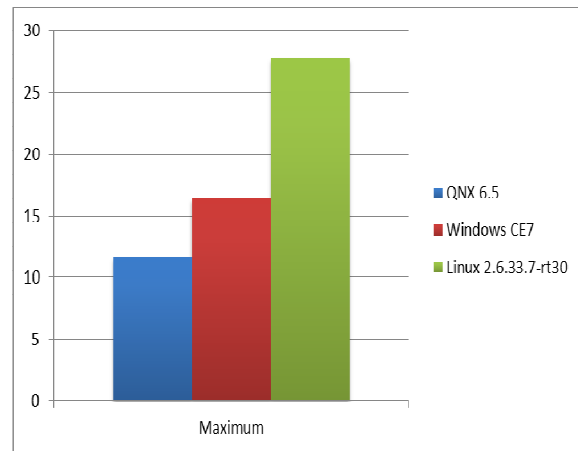
## 3.5.2 Thread switch latency between same priority threads (THR-P-SLS)

The “latency between threads of same priority” test measures the time to switch between threads of the same priority using SCHED\_FIFO policy. This test was performed four times, and each time using an increasing number of threads (2, 10, 128, and 1000) in order to generate the worst case behaviour.

The figures below present the thread switch latency with 1000 active threads in order to show the time values in the worst case. The evaluation results with fewer threads are presented in the next chapter.



**Figure 4a: Average latency between 1000 threads**



**Figure 4b: Maximum latency between 1000 threads**

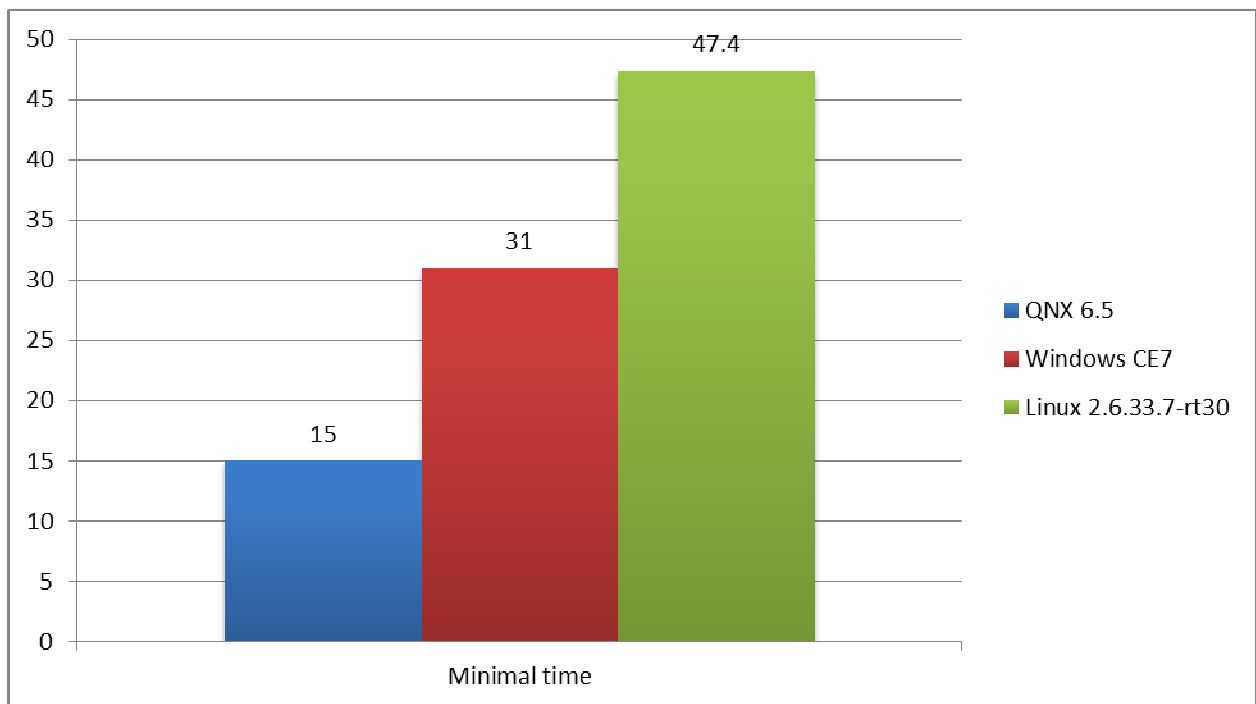
In this test, QNX Neutrino 6.5 outperforms Windows Embedded Compact 7 for average latency. Linux RT shows better results than Windows CE7.

The maximum thread latency depends on the clock tick interrupts; this means that if the test would run long enough, we would see similar results as the clock tick duration test (All tests show this congruence, which is why the clock tick duration test is so important for evaluating a RTOS). For Linux RT, the maximum latency is somewhat greater than for QNX Neutrino and Windows CE7.

## 3.5.3 Maximum sustained interrupt frequency (IRQ\_S\_SUS)

The “maximum sustained interrupt frequency” test measures the probability that an interrupt might be missed. In other words, it attempts to answer the question: Is the interrupt handling duration stable and predictable?

In this test, 1 billion interrupts are generated at specific interval rates. Our test suite measures whether the system under test misses any of the generated interrupts. The test is repeated with smaller and smaller intervals until the system under test is deemed to no longer handle the interrupt load.



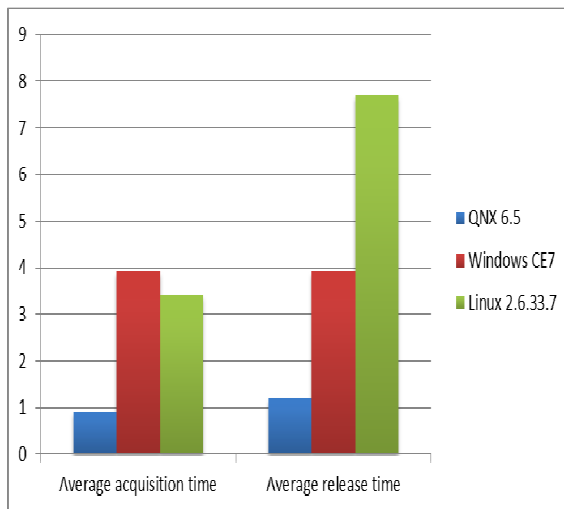
**Figure 5: The minimal interrupt period required in order not to lose any of the 1 billion interrupts**

QNX Neutrino fared best in handling the interrupts by successfully servicing interrupts generated every 15μs. Windows CE7 was second by handling interrupts with a value of 31 μs, while Linux RT functioned properly as long as interrupt levels were 47.4μs or greater.

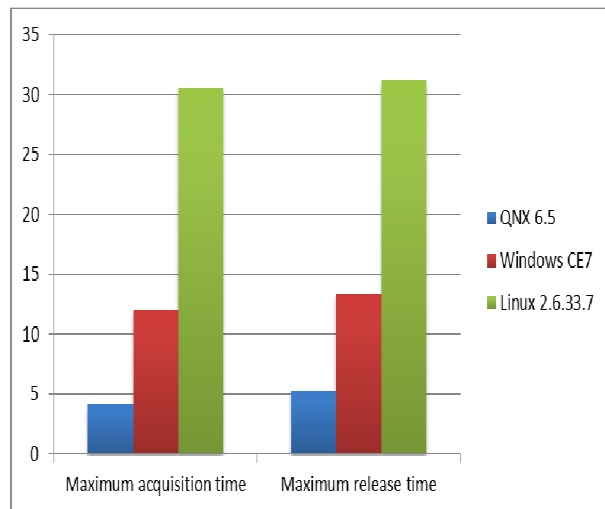


## 3.5.4 Mutex acquire-release timings: contention case (MUT-P-ARC)

The “mutex acquire-release timings in the contention case” test measures the time needed to acquire and release a mutex using priority inheritance. The acquire time is measured from the moment the higher priority thread requests the mutex until the moment the lower priority thread owning the mutex activates. The release time is measured from the moment the lower priority thread releases the mutex until the moment the higher priority thread is activated. As a result, the total time spent on a locked mutex is thus the sum of the acquisition time + release time + the time the lock is taken by the lower priority thread.



**Figure 6a: Mutex average acquire-release time: contention case**




**Figure 6b: Mutex maximum acquire-release time: contention case**

The advantage of the classical RTOSs compared with Linux is clear again. We also noticed that the release time on Linux RT is longer than expected. This is probably caused by the priority inheritance mechanism which takes some overhead, but it is required for real-time behavior.

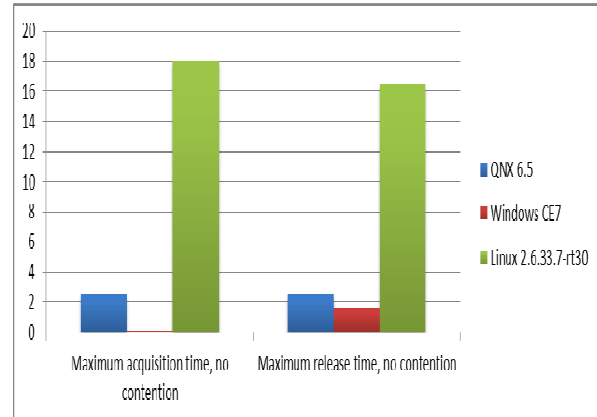
As the thread switch latency is on average larger on CE7 than on QNX, this difference shows up here as well (as the measurement includes a context switch).

## 3.5.5 Mutex acquire-release timings: no-contention case (MUT-P-ARN)

The “mutex acquire-release timings: no-contention case” test measures the overhead incurred using a lock when a thread is not locked by another thread.



**Figure 7a: Mutex average acquire-release time: no-contention case**



**Figure 7b: Mutex maximum acquire-release time: no-contention case**

In this case, the average is important as it shows how much time the locking overhead is in multi-threaded designs. Clearly, all OS tested here behave well. The maximum is in this case not that important as it shows us only the timer tick duration.

## 4 Detailed Comparison

This section presents the detailed test results and the comparison between the evaluated OSs.

### 4.1 Clock tests (CLK)

“Clock tests” measure the time that an operating system requires to handle its clock interrupts. On the tested platform, the clock tick interrupt is set on the highest hardware interrupt level, interrupting any other thread or interrupt handler.

#### 4.1.1 Clock tick processing duration (CLK-P-DUR)

The “clock tick processing duration” test examines the clock tick processing duration in the OS kernel. The test results are extremely important, as the clock interrupt will affect all the other performed measurements. The table below shows the average and maximum clock interrupt duration for the three tested OSs.

Clock interrupt duration	Average	Maximum
QNX Neutrino 6.5	2.2 $\mu$ s	10 $\mu$ s
Windows Embedded Compact 7	3 $\mu$ s	12 $\mu$ s
Linux RT 2.6.33.7	22 $\mu$ s	30 $\mu$ s




Figure 8a: Average clock interrupt duration




Figure 8b: Maximum clock interrupt duration

The clock tick processing time is important because it impacts latencies everywhere in the system. As we are interested in real-time behavior, the measurements for maximum processing times are more important than the measurements for average processing times. Our testing showed that the traditional RTOSs (QNX Neutrino and Windows CE) perform far better than Linux RT.

## 4.2 Thread tests (THR)

“Thread tests” measure the scheduler performance.

### 4.2.1 Thread creation behaviour (THR-B-NEW)

The “thread creation behavior” test examines the OS behavior when it creates threads. This test attempts to answer the question: Does the OS behave as it should in order to be considered a real-time operating system?

The following scenarios were checked in the test:

- If a thread is created with a lower priority than the creating thread, can we be sure that it will not be activated until the creating thread is finished?
- If a thread is created with the same priority as the creating thread, is it placed at the end of the ready queue?
- When yielding after it was created by a thread of the same priority (as in the previous scenario), does the newly created thread becomes active?
- If a thread is created with a higher priority than the creating thread, does this new thread become activate immediately?

<b>QNX Neutrino 6.5.0</b>	<b>Windows Embedded Compact 7</b>	<b>Linux 2.6.33.7-rt30</b>
Successfully passed this test	Successfully passed this test	Successfully passed this test

**Table 6: Results for the thread creation test**

QNX Neutrino and Windows CE7 passed this test successfully without any problems.



However, in Linux we observed different behaviors depending on whether SCHED\_FIFO or the SCHED\_RR class was used. When lowering the priority of a thread, then this thread:

- is placed at the head of the ready queue if the Linux OS is running with SCHED\_RR policy
- is placed at the end of the ready queue if the Linux OS is running with SCHED\_FIFO policy

Note that changing priorities at run-time is equivalent to dynamically creating and deleting threads, something that should not be done in a real-time system.

## 4.2.2 Round robin behaviour (THR-B-RR)

The “round robin behavior” test checks if the scheduler uses a fair round robin mechanism to schedule threads that: use the SCHED\_RR scheduling policy, are of the same priority, and are in the ready-to-run state (and using)!

<b>QNX Neutrino 6.5.0</b>	<b>Windows Embedded Compact 7</b>	<b>Linux 2.6.33.7-rt30</b>
Successfully passed this test	Passed 	Passed 

**Table 7: Results of the round robin test**

 Note that:

- For the **Linux** scheduler, the initial time slice of a created thread is 10 times greater than other slices (1second instead of the default 100milliseconds (ms)).
- For **Windows Embedded Compact 7**: the initial time slice of a created thread is 10 times greater than other slices (100 milliseconds (ms) instead of 10 milliseconds where 1ms = 1tick). By default threads have a 100ms time slice upon creation. When changing this default value, it will only be taken into account after the first time slice passed.

Since dynamic thread creation and the use of different threads with the same priority is a poor practice in real-time systems, we assumed that real-time projects would avoid these practices. Given this assumption, we discounted these issues when we determined the final scores of the OSs we evaluated.

## 4.2.3 Thread switch latency between same priority threads (THR-P-SLS)

The “thread switch latency between same priority threads” test measures the time needed to switch between threads of the same priority. For this test, threads must voluntarily yield the processor for other threads.

In this test, we use the SCHED\_FIFO policy. If we do not use the “first in first out” policy, a round-robin clock event could occur between the yield and the trace, so that the thread activation is not seen in the trace.

This test was performed in order to generate the worst-case behavior. We performed the test with an increasing number of threads, starting with two (2) and going up to 1000 in order to observe the behavior in a worst-case scenario. As we increase the number of active threads, the caching effect becomes evident since the thread context will no longer be able to reside in the cache.

Test	QNX		Windows CE7		Linux RT	
	Avg	Max	Avg	Max	Avg	Max
Thread switch latency, 2 threads	0.4	2.8	2.1	5.9	0.6	27.7
Thread switch latency, 10 threads	0.4	11.9	2.3	6.7	0.6	24.2
Thread switch latency, 128 threads	0.6	11.5	3.7	14.6	1.7	34.6
Thread switch latency, 1000 threads	0.8	11.7	4.4	16.4	2	27.8

Table 8: Thread switch latency between x threads, in  $\mu$ s




Figure 9a: Average switch latency between x threads, in µs




Figure 9b: Maximum switch latency between x threads, in µs

The impact of the caches on the average results is clearly observable (Figure 9a): the more threads there are to switch between, the more there are cache misses. The maximum values (Figure 9b) depend largely on the clock tick duration.

Context switching between threads is something that could be improved in CE7.



## 4.2.4 Thread creation and deletion time (THR-P-NEW)


The “thread creation and deletion time” test examines the time required to create a thread, and the time required to delete a thread in the following different scenarios:

- **Scenario 1 “never run”:** The created thread has a lower priority than the creating thread and is deleted before it has any chance to run. No thread switch occurs in this test.
- **Scenario 2 “run and terminate”:** The created thread has a higher priority than the creating thread and will be activated. The created thread immediately terminates itself (thread does nothing).
- **Scenario 3 “run and block”:** The same as the previous scenario (scenario 2: run and terminate), but the created thread does not terminate; it lowers its priority when it is activated, and therefore blocks.

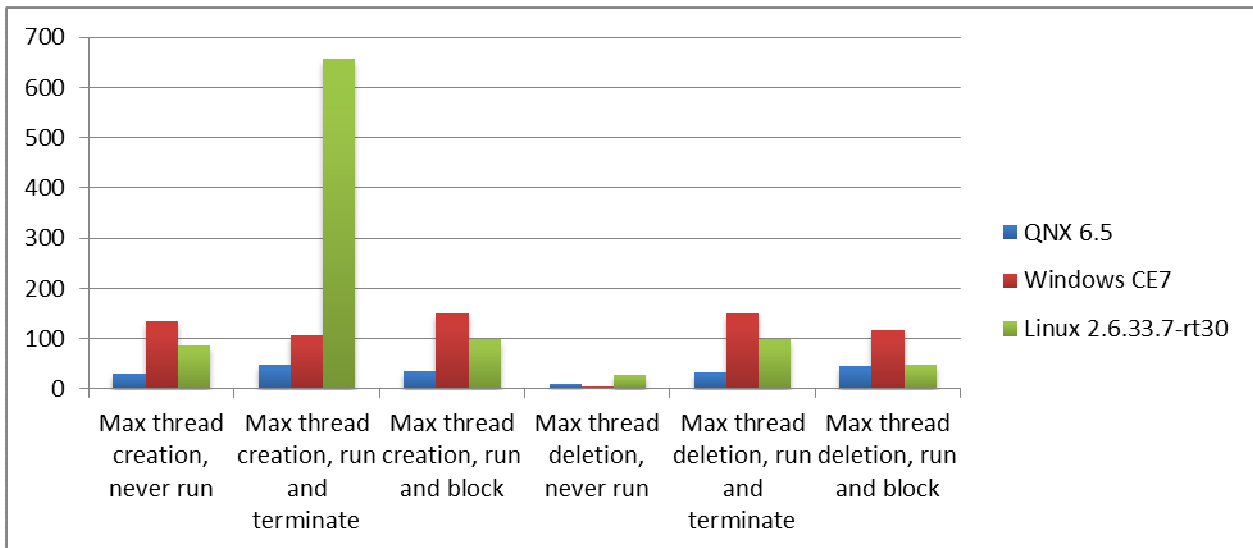
In the scenarios where the thread actually runs (2, 3), the creation time is the duration elapsed between the system call creating the thread and the moment the created thread is activated. For the “never run” scenario, the creation time is the duration of the system call.

Test		QNX		Windows CE7		Linux RT	
		Avg	Max	Avg	Max	Avg	Max
Never run	Thread creation	24	30.3	51.6	136	25.3	85.7
	Thread deletion	18.5	49	45.6	106	34.8	656.7
Run and terminate	Thread creation	23.8	34.5	56.2	149	38.1	100.2
	Thread deletion	1.9	9	2.1	5.4	0.6	28.2
Run and block	Thread creation	23.9	31.9	56.1	149	38.2	100.4
	Thread deletion	18.6	47	46.6	117	20.3	48.9

**Table 9: Thread creation and deletion testing results, in µs**



**Figure 10a: Average thread creation and deletion times (µs) in different scenarios**



**Figure 10b: Maximum thread creation and deletion times (µs) in different scenarios**

As dynamic threads is a bad design principle in real-time systems, the measurements here are more informative and will not impact the real-time performance score.

## 4.3 Semaphore tests (SEM)

“Semaphore tests” examine the behavior and performance of the OS counting semaphore. The counting semaphore is a system object that can be used to synchronize threads.

With all the operating systems we tested, we did not specify a name to the semaphore when we conduct our tests. An unnamed semaphore cannot be used between processes. This limitation does not necessarily mean that the implementation with an unnamed semaphore does not use round-trips to the kernel.

### 4.3.1 Semaphore locking test mechanism (SEM-B-LCK)

In this test, we verify if the counting semaphore locking mechanism works as it is expected to work. If this mechanism works as expected, then:

- The P ( ) call will block only when the count is zero.
- The V ( ) call will increment the semaphore counter.
- In the case where the semaphore counter is zero, the V ( ) call will cause a rescheduling by the OS, and blocked threads may become active.

QNX Neutrino 6.5.0	Windows Embedded Compact 7	Linux 2.6.33.7-rt30
The semaphore behaves correctly as a protection mechanism	The semaphore behaves correctly as a protection mechanism	The semaphore behaves correctly as a protection mechanism

**Table 10: Semaphore behaviour results**

## 4.3.2 Semaphore releasing mechanism (SEM-B-REL)

The “semaphore releasing mechanism” test verifies that the highest priority thread being blocked on a semaphore will be released by the release operation. This action should be independent of the order of the acquisitions taking place.

<b>QNX Neutrino 6.5.0</b>	<b>Windows Embedded Compact 7</b>	<b>Linux 2.6.33.7-rt30</b>
Successfully passed this test	Successfully passed this test	Successfully passed this test

**Table 11: Semaphore release mechanism testing results**

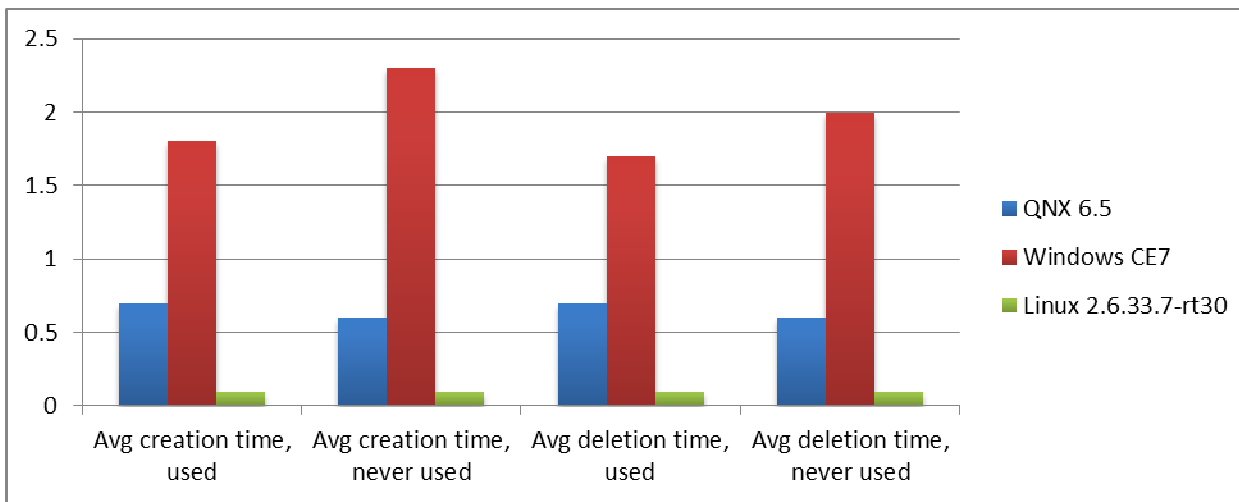
## 4.3.3 Time needed to create and delete a semaphore (SEM-P-NEW)

The “time needed to create and delete a semaphore” test measures the time needed to create a semaphore and the time needed to delete it. The deletion time is checked in two cases:

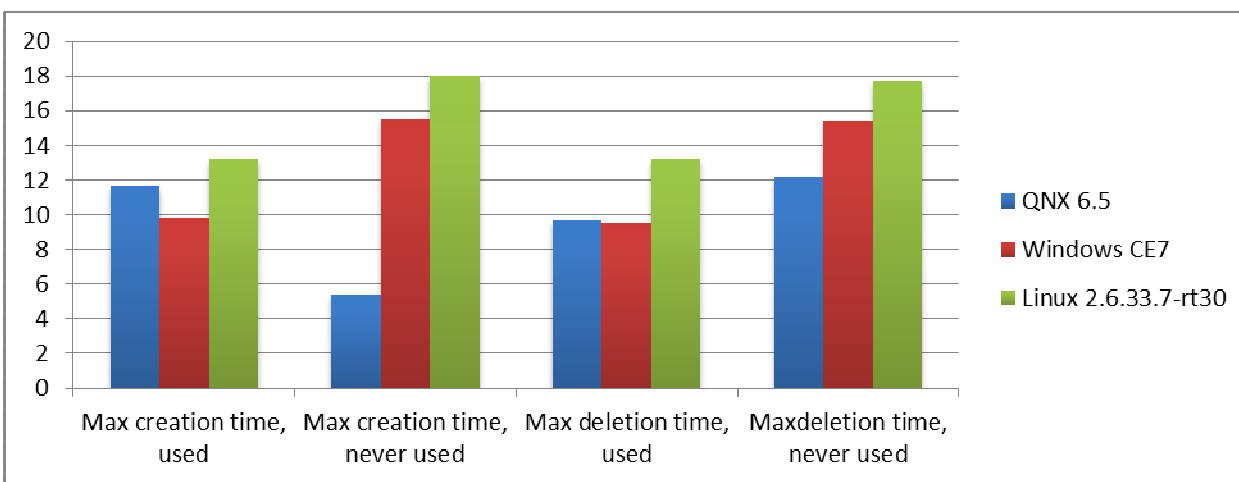
- The semaphore is used between the creation and deletion.
- The semaphore is not used between the creation and deletion.

Test		QNX		Windows CE7		Linux RT	
		Avg	Max	Avg	Max	Avg	Max
Semaphore is used	Creation time	0.7	11.7	1.8	9.8	<0.1	13.2
	Deletion time	0.6	5.4	2.3	15.5	<0.1	18
Semaphore is never used	Creation time	0.7	9.7	1.7	9.5	<0.1	13.2
	Deletion time	0.6	12.2	2	15.4	<0.1	17.7

**Table 12: Semaphore creation and deletion time results**



**Figure 11a: Average creating and deleting times (µs) of a Semaphore in different cases**



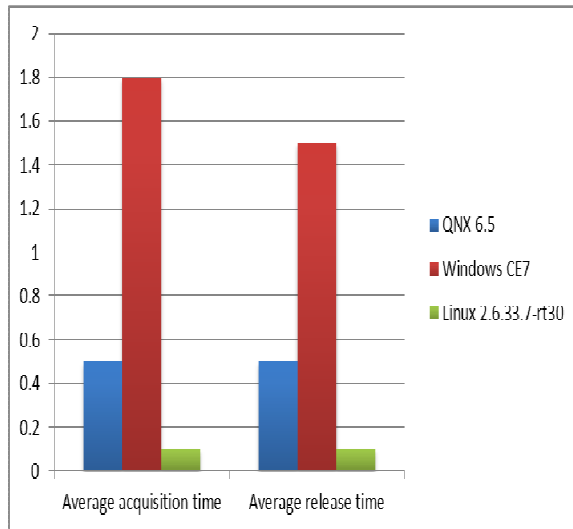
**Figure 11b: Maximum creating and deleting times (µs) of a Semaphore in different cases**

## 4.3.4 Test acquire-release timings: non-contention case (SEM-P-ARN)

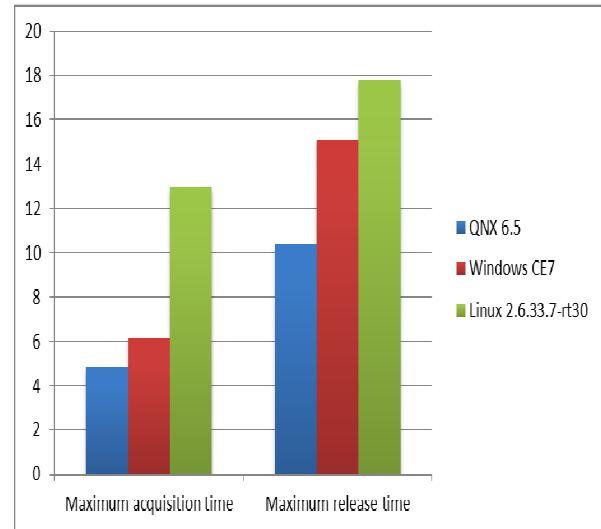
The “acquire-release timings: non-contention case” test measures the acquisition and release time in the non-contention case. Since in this test the semaphore does not neither block nor causes any rescheduling (thread switching), the duration of the call should be short.

Test	QNX		Windows CE7		Linux RT	
	Avg	Max	Avg	Max	Avg	Max
Semaphore acquisition time, no contention	0.5	4.8	1.8	6.1	0.1	13
Semaphore release time, no contention	0.5	10.4	1.5	15.1	0.1	17.8

**Table 13: Acquire release timings in the non-contention case**



**Figure 12a: Semaphore average acquire-release time: no contention**



**Figure 12b: Semaphore maximum acquire-release time: no contention**

Both QNX and CE7 perform a round trip to the kernel in these cases (like that required for named semaphores) while Linux uses atomic instructions and does not need a round-trip.

Note, however, that semaphores are much more appropriate way for signaling threads than for use as a protection mechanism. Indeed, semaphores do not have the concept of ownership and thus cannot be used to prevent priority inversions. If protection is required, mutexes should be used instead.

## 4.3.5 Test acquire-release timings: contention case (SEM-P-ARC)

The “acquire release timings: contention case” test is performed to test the time needed to acquire and release a semaphore, depending on the number of threads blocked on the semaphore. It measures the time in the contention case when the acquisition and release system call causes a rescheduling to occur.

The purpose of this test is to see if the number of blocked threads has an impact on the times needed to acquire and release a semaphore. It attempts to answer the question: “How much time does the OS needs to find out which thread should be scheduled first?”

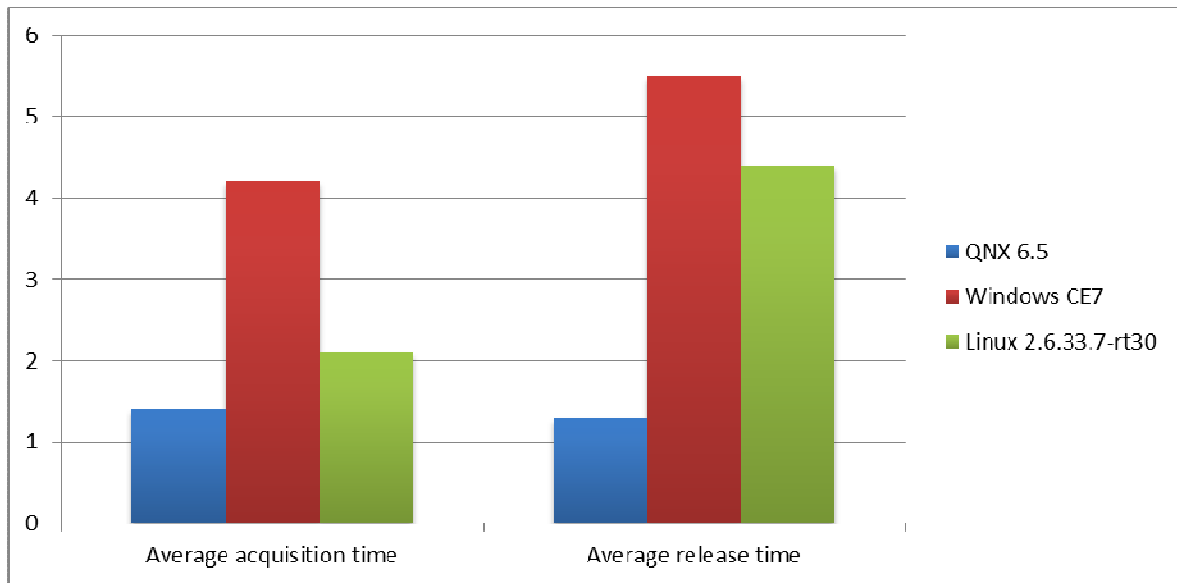
In this test, since each thread has a different priority, the question is how the OS handles these pending thread priorities on a semaphore. For more precise understanding of our test, please see the expanded diagrams showing a small time frame (e.g. one test loop). These diagrams are found in [Doc 5] for QNX Neutrino, [Doc 7] for Linux RT, and [Doc 9] for Windows Embedded Compact 7.

The test is conducted as follows:

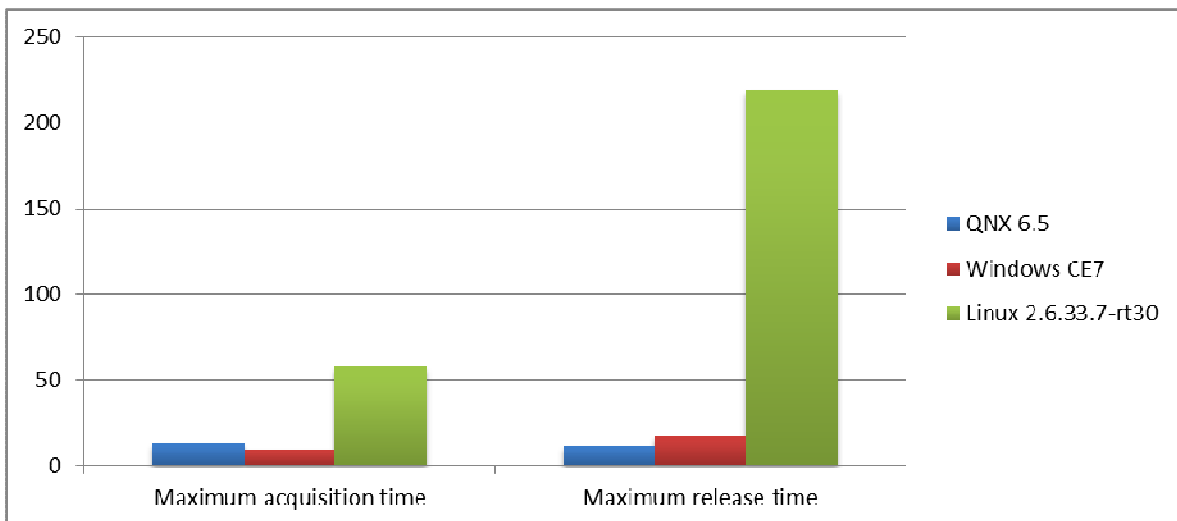
- We create a semaphore with count zero: so it will block on acquire.
- We create 128 threads with different priorities. The creating thread has a lower priority than the threads being created.
- When a thread starts execution, it tries to acquire the semaphore; but as the semaphore is taken, the thread stops and the kernel switches back to the creating thread. The time from the acquisition attempt (which fails) to the moment the creating thread is activated again, is called here the “acquisition time”. This time includes the thread switch time.
- Thread creation takes some time; so the time between each measurement point is large compared with most other tests.
- After the last thread is created and is blocked on the semaphore, the creating thread starts to release the semaphore, repeating this action the same number of times as the number of blocked threads on the semaphore.
- We start timing from the moment the semaphore is released, which in turn activates the pending thread with the highest priority, which stops the timing. Again, the thread switch time is included in the measurement.

Test	QNX		Windows CE7		Linux RT	
	Avg	Max	Avg	Max	Avg	Max
Semaphore <b>acquisition</b> time, contention	1.4	12.8	4.2	9.3	2.1	57.6
Semaphore <b>release</b> time, contention	1.3	12	5.5	17.3	4.4	219





**Figure 13a: Semaphore average acquire-release time: Contention**



**Figure 13b: Semaphore maximum acquire-release time: Contention**

Linux RT exhibited a strange behaviour: we noticed that the clock tick interrupt duration time increases on release, influencing seriously the worst case behaviour!

Windows Embedded Compact 7 thread switch overhead can be seen here as well.

## 4.4 *Mutex tests (MUT)*

Our “mutex tests” help us evaluate the behavior and performance of the mutual exclusive semaphore.

Although the mutual exclusive semaphore (further called mutex) is usually described as being the same as a counting semaphore where the count is one, this is not true. The behavior of a mutex is completely different than the behavior of a semaphore. Unlike semaphores, mutexes use the concept of a “lock owner”, and can thus be used to prevent priority inversions. Semaphores cannot do this, and it goes without saying that mutexes (and not semaphores) should not be used semaphores for critical section protection mechanisms. In scope of the framework, this test will look into detail of a mutex system object that avoids priority inversion.

Our test will on purpose generate a priority inversion with three threads:

- Low priority thread having a lock
- Intermediate priority thread ready to run
- High priority thread running and requesting the lock owned by the low priority thread

If the mutex has some priority inversion avoidance mechanism present, the intermediate priority thread may not run until the lower priority thread released the mutex and the high priority thread finished its work.

Without such avoidance mechanism, the intermediate priority thread will start to run and thus delay the higher priority thread. Thus, as a result, priorities would be inverted!

### 4.4.1 *Priority inversion avoidance mechanism (MUT-B-ARC)*

The “priority inversion avoidance mechanism” test determines if the system call being tested prevents the priority inversion case. To check this possibility, the test artificially creates a priority inversion.

<b>QNX Neutrino 6.5.0</b>	<b>Windows Embedded Compact 7</b>	<b>Linux 2.6.33.7-rt30</b>
Priority inversion is prevented as expected	Priority inversion is prevented as expected	Priority inversion is prevented as expected

**Table 14: results of priority inversion avoidance mechanism**

## 4.4.2 Mutex acquire-release timings: no-contention case (MUT-P-ARN)

The “mutex acquire-release timings: no contention case” test measures the overhead incurred by using a lock when this lock is not owned by any other thread. Well-designed software will use non-contended locks most of the time, and only in some rare cases the lock will be taken by another thread.

Therefore, it is important that the non-contention case should be fast. Note that the required speed is only possible if the CPU supports some type of atomic instruction, so that no system call is needed when no contention is detected.

Test	QNX		Windows CE7		Linux RT	
	Avg	Max	Avg	Max	Avg	Max
Mutex <b>acquisition</b> time, no-contention	0.1	2.5	<0.1	<0.1	0.1	18.1
Mutex <b>release</b> time, no-contention	0.1	2.5	<0.1	1.6	0.1	16.5

Table 15: Results of the mutex acquire-release timing in no-contention case, in  $\mu$ s

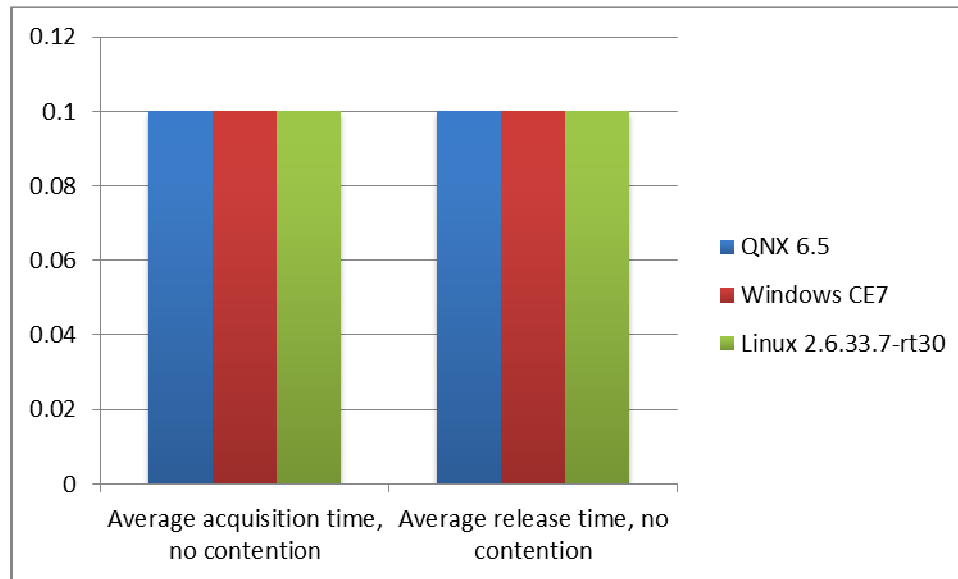
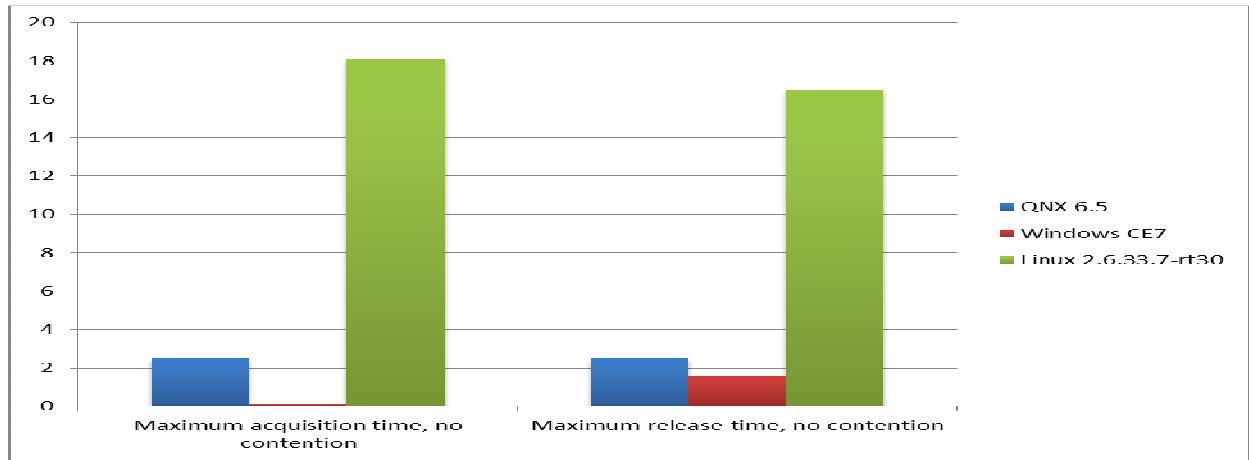


Figure 14a: Mutex average acquire-release time: no-contention



**Figure 14b: Mutex maximum acquire-release time: no-contention**

The average acquire-release time differences are too small to be measured. It means that all OSes tested here handle this well. The aim is indeed to have a minimal impact when using locks which, in a good design, will mostly be used without contention.

However, the maximum values were measureable: Figure 14b presents the clock tick durations, in microseconds ( $\mu$ s).

### 4.4.3 *Mutex acquire-release timings: contention case (MUT-P-ARC)*

The “mutex acquire-release timings: contention case” test is the same test as the “priority inversion avoidance mechanism (MUT B ARC)” test described above, but performed in a loop. In this case, we measure the time needed to acquire and release the mutex in the priority inversion case.

Our test is designed so that the acquisition enforces a thread switch:

- The acquiring thread is blocked
- The thread with the lock is released.

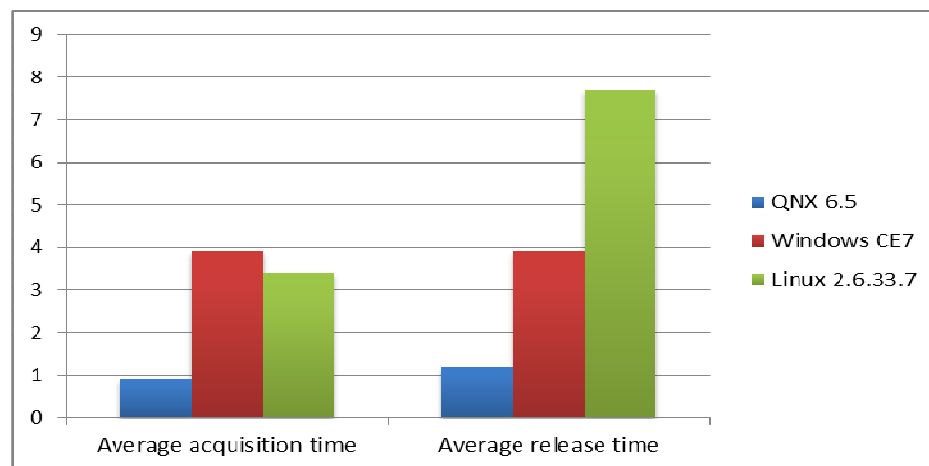
We measured the acquisition time from the request for the mutex acquisition to the activation of the lower priority thread with the lock.

Note that before the release, an intermediate priority level thread is activated (between the low priority one having the lock and the high priority one asking the lock). Due to the priority inheritance, this thread does not start to run (the low priority thread having the lock inherited the high priority of the thread asking the lock).


We measured the release time from the release call to the moment the thread requesting the mutex was activated; so this measurement also includes a thread switch.

Test	QNX		Windows CE7		Linux RT	
	Avg	Max	Avg	Max	Avg	Max
Mutex <b>acquisition</b> time, contention	<b>0.9</b>	<b>4.2</b>	<b>3.9</b>	<b>12</b>	<b>3.4</b>	<b>30.5</b>
Mutex <b>release</b> time, contention	<b>1.2</b>	<b>5.3</b>	<b>3.9</b>	<b>13.3</b>	<b>7.7</b>	<b>31.2</b>

**Table 16: Results of the mutex acquire-release timing in contention case, in  $\mu$ s**



**Figure 15a: Mutex average acquire-release time: contention**



**Figure 15b: Mutex maximum acquire-release time: contention**

The Windows Embedded Compact 7 larger thread context switch overhead is seen here again.

## 4.5 Interrupt tests (IRQ)

“Interrupt tests” evaluate how the operating system performs when handling interrupts.

Interrupt handling is a key system capability of real-time operating systems. Indeed, RTOSs are typically event driven.

For our interrupt tests,, our standard tracing system is adapted. Interrupts are generated by a plugged-in PCI related card (can be PMC/PCI or CPCI). This card has a complete independent processor on board, with custom-made software. As such, we can guarantee that the independent interrupt source is not synchronised in any way with the platform under test.

### 4.5.1 Interrupt latency (IRQ\_P\_LAT)

The “interrupt latency” test measures the time it takes to switch from a running thread to an interrupt handler. This time is measured from the moment the running thread is interrupted, so the measurement does not take into account the hardware interrupt latency.

Test	QNX		Windows CE7		Linux RT	
	Avg	Max	Avg	Max	Avg	Max
Interrupt latency	1.7	2.8	6.7	16.7	10.4	33.4

Table 17: Interrupt latency results in  $\mu$ s




Figure 16a: Interrupt average latency




Figure 16b: Interrupt maximum latency

For average interrupt latency, QNX Neutrino 6.5 showed very good results. Windows Embedded Compact 7 is already a bit slower (due to the slower context switch). Linux RT is still a bit slower, even if it has shorter context switch latency.

However, for real-time systems, average interrupt latency is less important than maximum latency, or worst case. Of course, there is no easy way to set an upper limit of the worst case. Using a statistical approach, you just need more samples to have a more accurate view. That's the reason why we run a long duration interrupt test. Hence, the maximum sustained interrupt rate is the most important test (section 4.5.4) for the real-time behavior.



## 4.5.2 Interrupt dispatch latency (IRQ\_P\_DLT)

The “interrupt dispatch latency” test measures the time the OS takes to switch from the interrupt handler back to the interrupted thread.

Test	QNX		Windows CE7		Linux RT	
	Avg	Max	Avg	Max	Avg	Max
Dispatch latency from interrupt handler	1.2	3.4	3	8.5	11.5	60

Table 18: Interrupt dispatch latency in  $\mu$ s





Figure 17a: Average dispatch latency from interrupt handler    Figure 17b18: Maximum dispatch latency from interrupt handler

As discussed before, Windows Embedded Compact 7 uses an interrupt thread. So handling an interrupt involves a thread switch.

Note that the results for the maximum dispatch latency depend on whether or not the test catches timer interrupts.

## 4.5.3 Interrupt to thread latency (IRQ\_P\_TLT)

The “interrupt to thread latency” test measures the time the OS takes to switch from the interrupt handler to the thread that is activated from the interrupt handler.

Not all the evaluated OSs handle switching in the same way, and we tailored our tests to obtain comparable results:

- For QNX Neutrino, the interrupt handler emits an event to release a blocked thread. This blocked thread has the highest priority in the system.
- For Linux RT, a thread is blocked by using an `ioctl` call, and is released in the kernel module upon the interrupt.
- For Windows CE7, the thread uses a `WaitForSingleObject` on an interrupt event.

This test measures the time the OS takes from the interrupt handler to the blocked thread (as a consequence this includes a thread switch).

Test	QNX		Windows CE7		Linux RT	
	Avg	Max	Avg	Max	Avg	Max
Latency from ISR to waken-up thread	3.2	14.7	2.7	8.8	4.9	29.9

Table 19: Interrupt to thread latency, in  $\mu$ s




Figure 18a: Average latency from ISR to waken-up thread, in  $\mu$ s




Figure 18b: Maximum latency from ISR to waken-up thread, in  $\mu$ s


## 4.5.4 Maximum sustained interrupt frequency (IRQ\_S\_SUS)

The “maximum sustained interrupts frequency” test measures the probability that an interrupt might be missed. It attempts to answer the question: Is the interrupt handling duration stable and predictable?

In this test we load the system with a high load interrupt source which generates 1 billion interrupts and determine at which interrupts frequency the OS begins to miss interrupts. The table below shows the minimum delay required between interrupts for the OSs tested to not lose any of the 1 billion interrupts. Below this threshold, the OSs lost interrupts.

Note that this test presents the worst case of the best-case scenario: due to the high interrupt rate, the interrupt handler is expected to be in the cache all time. Nevertheless, we observed clear differences in the performance of the OSs we tested. In order to not miss any interrupts, Linux (with the RT patch) requires around three (3) times more time between interrupts than does QNX Neutrino 6.5. Windows CE7 requires 31  $\mu$ s which is almost twice the value of QNX Neutrino who had the best time (15  $\mu$ s) in handling interrupts.

Test	QNX	Windows CE7	Linux RT
Minimal interrupt period required not to lose any of the <u>1 billion</u> generated interrupts.	15	31	47.4



**Figure 19: Minimal interrupt period required not to lose any of the 1 billion generated interrupts.**

## 5 Conclusion

Our first conclusion is that both traditional RTOSs QNX Neutrino and Windows CE7 can be qualified as true real time operating systems, out-of-the-box.

Linux with the RT\_PREEMPT patch can also be qualified as RTOS, though for this OS the user must take care to use a correct kernel configuration, both at build time and at run time.

However, the behavior and performance gap that separated Linux RT from both QNX Neutrino and Windows CE7 still does not make Linux RT a comparable contender as a real-time OS.

## 6 Related documents

These are documents that are closely related to this document. They can all be downloaded using following link: <http://download.dedicated-systems.com/>

- |        |  |
|--------|--|
| Doc. 1 | The evaluation framework<br>This document presents the evaluation framework. It also indicates which documents are available, and how their name giving, numbering and versioning are related. This document is the base document of the evaluation framework.<br>EVA-2.9-GEN-01                      Issue: 1                      Date: April 19, 2004 |
| Doc. 2 | The evaluation test report definition.<br>This document presents the different tests issued in this report together with the flowcharts and the generic pseudo code for each test. Test labels are all defined in this document.<br>EVA-2.9-GEN-03                      Issue: 1                      April 19, 2004                                     |
| Doc. 3 | The OS evaluation template<br>This document presents the layout used for all reports in a certain framework.<br>EVA-2.9-GEN-04                      Issue: 1                      April 19, 2004   |
| Doc. 4 | QNX v6.5, Theoretical evaluation report<br>This document presents the qualitative discussion of the QNX OS<br>EVA-2.9-OS-QNX-65                      Issue: 4.1                      Sept 8, 2011  |
| Doc.5  | QNX technical evaluation report<br>This document presents the results of evaluating QNX on ARM platform<br>EVA-2_9-TST-QNX-65-x86A-02    Issue: 3.2                      Sept 7, 2011  |
| Doc. 6 | Linux theoretical evaluation report<br>This document presents the qualitative discussion of the Linux OS<br>EVA-2.9-OS-LINUXRT_2.6.33.7.2-rt30    Issue: 1.07                      May 30, 2011  |
| Doc.7  | Linux technical evaluation report<br>This document presents the results of evaluating Linux on ARM platform<br>EVA-2_9-TST-LINUXRT_2_6_33_7_2-rt30-ATOM    Issue: 2    Apr 26, 2012  |
| Doc. 8 | Windows Embedded Compact 7 theoretical evaluation report<br>This document presents the qualitative discussion of the Windows Embedded Compact 7 OS<br>EVA-2.9-OS-CE-7-A03                      Issue: 2.1                      Sept 19, 2011   |
| Doc.9  | Windows Embedded Compact 7 technical evaluation report<br>This document presents the results of evaluating Windows Embedded Compact 7 on ARM platform<br>EVA-2.9-TST-CE7-ATOM                      Issue: 2                      May 1, 2012   |

## 7 Appendix A: Acronyms

Acronym	Explanation
API	Application Programmers Interface: calls used to call code from a library or system.
BSP	Board Support Package: all code and device drivers to get the OS running on a certain board
DSP	Digital Signal Processor
FIFO	First In First Out: a queuing rule
GPOS	General Purpose Operating System
GUI	Graphical User Interface
IDE	Integrated Development Environment (GUI tool used to develop and debug applications)
IRQ	Interrupt Request
ISR	Interrupt Servicing Routine
MMU	Memory Management Unit
OS	Operating System
PCI	Peripheral Component Interconnect: bus to connect devices, used in all PCs!
PIC	Programmable Interrupt Controller
PMC	PCI Mezzanine Card
PrPMC	Processor PMC: a PMC with the processor
RTOS	Real-Time Operating System
SDK	Software Development Kit
SoC	System on a Chip

Doc: **EVA-2.9-CMP-ATOM**

Issue: **v 1.00**

Date: **May 1, 2012**

## DOCUMENT CHANGE LOG

Issue No.	Revised Issue Date	Para's / Pages Affected	Reason for Change
0	Apr 18, 2012	ALL	Initial report
1	April 27, 2012	ALL	Final Report